Correlation-Based Pattern Recognition for Implantable Defibrillators

Jeff Wilkins, Ph.D.
Information Systems Laboratory
Department of Electrical Engineering
Stanford University
Stanford, California

An estimated 300,000 Americans die each year from cardiac arrhythmias. Historically, drug therapy or surgery were the only treatment options available for patients suffering from arrhythmias. Recently, implantable arrhythmia management devices have been developed. These devices allow abnormal cardiac rhythms to be sensed and corrected in vivo. Proper arrhythmia classification is critical to selecting the appropriate therapeutic intervention. The classification problem is made more challenging by the power/computation constraints imposed by the short battery life of implantable devices. Current devices utilize heart rate-based classification algorithms. Although easy to implement, rate-based approaches have unacceptably high error rates in distinguishing supraventricular tachycardia (SVT) from ventricular tachycardia (VT). Conventional morphology assessment techniques used in ECG analysis often require too much computation to be practical for implantable devices.

In this paper, a computationally-efficient, arrhythmia classification architecture using correlation-based morphology assessment is presented. The architecture classifies individual heart beats by assessing similarity between an incoming cardiac signal vector and a series of prestored class templates. A series of these beat classifications are used to make an overall rhythm assessment. The system makes use of several new results in the field of pattern recognition.

The resulting system achieved excellent accuracy in discriminating SVT and VT.

INTRODUCTION

An estimated 300,000 Americans die each year from ventricular arrhythmias[1]. Historically, once a patient was diagnosed with ventricular tachycardia (VT) or ventricular fibrillation (VF), treatment involving drug therapy or surgery was mandated. Both treatment options have serious shortcomings.

A third treatment option, the implantable cardioverter defibrillator (ICD)[2], has emerged over the last ten years offering radically improved outcomes for patients suffering from VT or VF. The ICD analyzes the electrical activity of the heart from electrodes attached directly to its surface. The sensed waveform is called an electrogram (EGM). When a ventricular tachyarrhythmia is sensed, pacing pulses or higher energy shocks are delivered to return the patient to a normal rhythm.

Since different therapeutic interventions are mandated for different arrhythmias, a key success factor for the ICD is error-free (or near error-free) arrhythmia classification.

Severe power constraints imposed on implantable devices complicate matters. Because of the relationship between dissipated power and computation speed, a power constraint effectively imposes a computation constraint. As a result, most existing ICDs make use of simple heart rate-based classification algorithms. Unfortunately, rate-based algorithms suffer notable shortcomings in discriminating supraventricular tachycardia (SVT) from ventricular tachycardia (VT).

Previous work has demonstrated that morphology assessment (e.g. measuring the shape of beats, instead of the time interval between them) can improve SVT/VT discrimination[2]. However, the computation required for quantitative morphology assessment has prevented its widespread use in ICDs.

In this paper minimum computation arrhythmia classification techniques using morphology assessment are developed. The proposed architecture (see Fig. 1) consists of a *preprocessor*, a *beat classifier* and a *rhythm classifier*. The preprocessor filters and digitizes the analog EGM. The beat classifier calculates the sample correlation coefficient between an EGM signal vector and a series of preconstructed VT and SVT beat templates. The signal vector is assigned to the beat class with whose template it has the maximum correlation coefficient. The beat classification is used as a building block for an SVT/VT rhythm classifier. The rhythm classifier assesses the more prevalent class amongst the five most recent beat classifications (majority decision rule) to determine the prevailing rhythm.

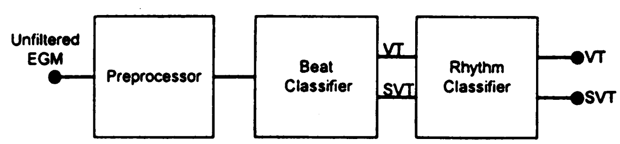

Figure 1 -System Block Diagram

In practice, the proposed architecture would augment a conventional rate-based approach.

METHODS

To be feasible, the proposed SVT/VT arrhythmia classification system must be both fast and accurate. Methods for optimal template construction, and acceleration techniques are first briefly presented. The reader is referred to the literature[3] for complete details. Finally, the SVT/VT discrimination of the proposed system is tested and compared to a rate-based system.

Beat Classifier Design Issues

Optimal Template Construction. A beat class template is constructed using a series of noisy realizations from that class. It can be shown[3] using the method of Lagrange multipliers that the optimal template (defined as that which maximizes the mean correlation coefficient between the template and a set of realizations) is formed by first normalizing each realization to be zero mean and unit length, then adding the resulting vectors together and renormalizing the result.

One potential complication is that the optimal template construction method assumes that the realizations are correctly registered (aligned). An iterative approach is required because the first, registration stage assumes that the template T is known, but this is not the case and estimation of T is the goal of the entire procedure. An initial guess is made of T, using the best initial estimate of the correct registration. T is subsequently refined through iteration (i.e. repeated registration and recomputation of T, until T is relatively stable).

The selection of the number of realizations to use in template construction is highly problem dependent. Using more beats requires greater computation but usually generates templates with higher signal-to-noise ratios. Five beats was experimentally determined to be adequate for the construction of good SVT/VT templates. Only two iterations were required.

Acceleration Techniques. Let S be a signal vector and T be a template vector, both of length N. Let $\bar{s}$ be the sample mean of the signal vector. Since the template is of zero mean and unit length, the product moment correlation coefficient is defined as:

$$r(S,T) \stackrel{\Delta}{=} \frac{\sum_{n=0}^{N-1}\left(s_n - \bar{s}\right)t_n}{\sqrt{\sum_{n=0}^{N-1}\left(s_n - \bar{s}\right)^2}} \qquad (1)$$

The computation of (1) for every possible registration of each template is very compute-intensive. Given that current generation ICDs make use of low power/low clock rate 8-bit microcontrollers (e.g. 100KHz 6502 in one leading model), significant acceleration must be obtained. Various acceleration methods are now described.

Template Size and Sampling Rate
The most obvious method of reducing the computational burden is to reduce the template size and sampling rate. Based upon a time and frequency domain analysis of SVT/VT beats, a 250 Hz sampling rate and 50-sample template was selected.

Registration-level Acceleration
Sum of Absolute Differences (SAD). A simple R wave detector may be used to determine, coarsely, when a beat is occuring. Unfortunately, such detectors can usually only determine the best alignment to within +/- 5 samples (40 msec). To perform morphology assessment, a more precise determination is required. Translational registration is the process of selecting the best alignment between the signal vector and a given template. Because multiplication is so computationally expensive on the CISC microcontrollers found in ICDs, it is undesirable to perform the registration with a correlation coefficient similarity metric. Registration is instead performed using *a sum of absolute differences* (SAD) metric.

To evaluate a given registration, SAD pointwise subtracts the template from the signal vector. The absolute value of the difference between template and signal samples are summed. The registration which produces the minimum SAD is deemed best. The search space for the optimal translational registration is restricted to +/-5 samples about the R-wave detector selected registration.

Registration Stopping. SAD may be accelerated using a technique called *registration stopping*. This technique makes use of the fact SAD is a non-decreasing function. Consider the calculation of the SAD for a given registration. If the partial SAD obtained after the first n samples for this registration exceeds the complete SAD value obtained for another registration, it cannot be the best registration. Under such circumstances, it is not necessary to complete the rest of the SAD calculation. In practice, a single checkpoint of n=15 produces the greatest accelerations.

Once the best registration has been found, it is necessary to calculate (1) for each template. Several acceleration techniques which speed the computation of correlation coefficient are now presented.

Template-level Acceleration
Template Stopping. To perform beat classification, we need only determine which template has the higher correlation coefficient with the signal vector. A version of the Triangle Inequality involving correlation coefficient may be developed. An acceleration technique called *template stopping* makes use of this inequality:

Given two templates T_i and T_j and a signal vector S to be classified,

$$If \quad r(S,T_i) > \sqrt{\frac{1+r(T_i,T_j)}{2}} \quad then\ r(S,T_i) > r(S,T_j)$$

The reader is referred to the literature[3] for a complete development and proof of the algorithm. Qualitatively, the template stopping technique states that if the correlation coefficient between S and the template, T_i exceeds a certain threshold (which is a function of $r(T_i, T_j)$) then T_i is guaranteed to be more correlated with the signal vector than T_j. Under these circumstances, we need not calculate the correlation coefficient between S and T_j.

Template Ordering. *Template ordering* determines the order in which the correlation coefficients are calculated. The general rule is to calculate the correlation coefficient of the most likely template(s) first.

In a non-stationary application like EGM beat classification, probabilities should be conditioned on recent events. A statistical likelihood predictor using recent beat classifications could be used to determine the likelihood of the next EGM beat. In the interest of computational simplicity, the last beat classification will be utilized to select the order of calculation.

Sample-level Acceleration

Sample Stopping. In calculating (1), it would be desirable to prematurely terminate the computation for obviously bad template candidates. Sample stopping is an acceleration technique that provides this capability. After the first k terms (where k is a problem dependent variable) are calculated, an upper and lower bound on the full correlation coefficient is calculated and used to assess whether the remaining terms should be computed.

For notational simplicity, let S and T be zero mean, and unit length. Divide the calculation of (1) into two pieces:

$$r(S,T) = \sum_{j=1}^{k} s_j t_j + \sum_{j=k+1}^{N} s_j t_j \tag{2}$$

By the Cauchy-Schwartz inequality, we may bound the second summation in (2) from above and below:

$$-\sqrt{\sum_{j=k+1}^{N} t_j^2}\sqrt{\sum_{j=k+1}^{N} s_j^2} \le \sum_{j=k+1}^{N} s_j t_j \le \sqrt{\sum_{j=k+1}^{N} t_j^2}\sqrt{\sum_{j=k+1}^{N} s_j^2} \tag{3}$$

Since we have normalized the signal vector to unit length, the inequality may be further simplified:

$$-\sqrt{\sum_{j=k+1}^{N} t_j^2} \le \sum_{j=k+1}^{N} s_j t_j \le \sqrt{\sum_{j=k+1}^{N} t_j^2} \tag{4}$$

Substituting (4) into (2) yields the desired upper and lower bound:

$$\sum_{j=1}^{k} s_j t_j - \sqrt{\sum_{j=k+1}^{N} t_j^2} \le r(S,T) \le \sum_{j=1}^{k} s_j t_j + \sqrt{\sum_{j=k+1}^{N} t_j^2} \tag{5}$$

Define:

$$r_{partial}(S,T) \stackrel{\Delta}{=} \sum_{j=1}^{k} s_j t_j \qquad \left\|T_{residual}\right\| \stackrel{\Delta}{=} \sqrt{\sum_{j=k+1}^{N} t_j^2} \tag{6}$$

Substituting into (5) yields:

$$r_{partial}(S,T) - \left\|T_{residual}\right\| \le r(S,T) \le r_{partial}(S,T) + \left\|T_{residual}\right\| \tag{7}$$

With the values of the template residual terms precomputed and stored, we may use the inequality to determine after calculating the first k terms of the correlation whether to complete or terminate the calculation. When combined with template stopping, sample stopping provides a potent acceleration in the two class case:

1) If the first k terms of the first template calculation are bounded from below by a value higher than the template stopping threshold, we need not finish the calculation. The first template is guaranteed to have a higher correlation coefficient with the signal vector than the second template. If this condition is not met, we must complete the calculation of the first template correlation coefficient.

2) Acceleration can also be achieved during the calculation of the correlation coefficient with the second template. If the upper bound for the second template correlation after k samples is less than the first template correlation, then there is no point in completing the computation of the second template correlation. The first template is guaranteed to have a higher correlation coefficient.

Conversely, if after k samples the lower bound for the second template correlation coefficient is greater than the first template correlation coefficient the second template is guaranteed to be a better match. Under these circumstances, we need not finish the calculation of the second correlation coefficient.

k can be selected based on its impact on the template residual. Typically, k is selected so as to bound r by a small, known residual. This has the desirable effect of enhancing the efficacy of sample stopping and template stopping. In practice, the selection of k is highly problem dependent.

Sample Ordering. A little thought should confirm that not all pairs of sample and template points are created 'equally'. Those pairs with the greatest, positive product contribute the most to the value of the whole summation. Calculating the terms of correlation coefficient out-of-order can reduce the residuals in sample stopping.

If the templates points are sorted in order of descending absolute value, the residual in (7) will be minimized. The computational burdens of sample ordering are minimal, since sorting is only done whenever the template is updated.

For EGM beat templates, k=15 usually provides a residual of less than 0.01.

Rhythm Classifier Design Issues
The number of beats to use for rhythm assessment is a design tradeoff between time to classification and probability of error. If the beat errors are assumed to be independent, then it is a simple matter to use the binomial formula to determine the number of beats required to achieve a desired rhythm classification error rate.

For example, with an individual beat error rate of 0.01, a five beat rhythm assessment should have an error rate of 10^{-5}.

A more detailed analysis, which includes the possibility of non-SVT or VT beats (such as PVCs) is included in Wilkins[3].

Preprocessor Design
The preprocessor consists of an anti-alias filter and an A/D converter.

To minimize computation, the sampling rate is set to the lowest level consistent with good performance. Most commonly an anti-alias filter is designed as a low pass filter with an upper cutoff selected to prevent aliasing.

Based upon a temporal and spectral analysis of the EGM, a 12-bit, 250 Hz A/D converter was chosen. The choice of the upper cutoff frequency for the anti-alias filter is somewhat constrained. Allowing for a physically-realizable transition band, the cutoff cannot exceed 115 Hz and still prevent aliasing. Conversely, the cutoff cannot be reduced substantially below 115 Hz or classification accuracy is compromised.

This simplistic approach to the design of an anti-alias filter masks a deeper question. Since filtering changes the morphology of beats, it can have a significant impact on the efficacy of the acceleration techniques. Filtering can thus rightly be thought of as an acceleration technique in its own right. What impact does filtering have on the computational requirements of the proposed architecture? What is the 'optimal' filter?

Because of the difficulty of realizing analog filters with arbitrary transfer functions, the search space was restricted to easily implemented bandpass designs. A 121-tap linear phase FIR bandpass filter was used to simulate the effects of an analog anti-alias filter. An optimal lower cutoff frequency of 10 Hz and upper cutoff of 115 Hz was determined using non-linear programming techniques. The 10-115Hz band pass filter reduced computation by nearly one-third versus a conventional 115 Hz low pass anti-alias filter.

Experimental Design
Clinical data was gathered to test the efficacy of the proposed system. SVT/VT discrimination was evaluated for the proposed architecture and compared to a rate-based approach.

Patient Population. Eighteen recordings from eight patients (6 men/2 women) aged 52 to 85 were obtained. In all, a test set (exclusive of the samples used to construct templates for each patient) of 387 five beat VT segments and 944 five beat SVT segments were obtained. Data were gathered pursuant to an IRB-approved study by Ventritex, Inc. Written consent was obtained from participants.

Data Acquisition. Bipolar (1 cm) distal ventricular electrograms were recorded during routine clinical studies in the cardiac electrophysiology laboratory (see Figure 2). One 8 French and two 6 French side-arm sheaths (Cordis Corp., Miami, FL, USA) were positioned in the right femoral vein using the Seldinger technique. Three 6 French quadrapolar electrode catheters (USCI Division, C.R. Bard Inc., Billerica, MA, USA) with an interelectrode distance of 1 cm were introduced and advanced under fluoroscopic guidance. One electrode catheter was positioned in the high right atrium or right atrial appendage. Two electrode catheters were positioned in the right ventricular apex for right ventricular apex pacing and recording, respectively.

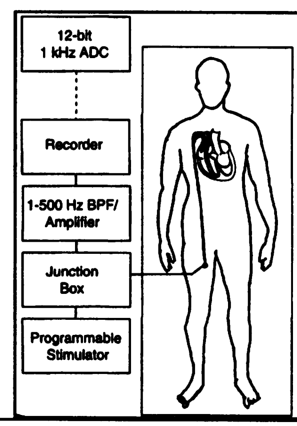

Figure 2 -Clinical setup for Data Acquisition

Sustained monomorphic ventricular tachycardia and supraventricular tachycardia were induced by programmed stimulation or alternating current. After signal amplification and bandpass filtering (1-500 Hz) bipolar intraventricular electrograms were recorded continuously on FM magnetic tape at a tape speed of 9.5 cm/sec (Hewlett-Packard Model 3968A, San Diego, CA, USA). The recorded ventricular electrograms were subsequently replayed and digitized on an IBM PC/XT personal computer at 1KHz with 12-bit resolution using a Tecmar Lab Master (Scientific Solutions, Inc., Solon, OH, USA) Analog-to-Digital Converter. Based upon subsequent spectral analysis, the data were decimated to a 250 Hz sampling rate using a 101 tap Chebyshev low pass filter with a 115 Hz cutoff frequency. The filter provides attention of more than 30 dB for frequencies above 125 Hz .

Expert Classification. Each recording was annotated by a cardiac electrophysiologist to ensure an accurate interpretation of each arrhythmia. Individual beats were classified as well as rhythm classifications of segments of five beats. To ensure proper classification, the electrophysiologist made use of atrial and ventricular EGM recordings as well as a surface ECG.

The expert was used as a gold standard to provide the 'correct' classification for each EGM beat and rhythm segment. The classification of the correlation-based architecture and a rate-based algorithm was then compared to that of the human expert.

Correlation-Based Classification. The correlation-based arrhythmia classification system described in this paper was used to perform SVT/VT rhythm classification.

Rate-Based Classification. A rate-based approach utilizing interval regularity was implemented to discriminate SVT and VT rhythms. Segments of five consecutive beats were analyzed. If the minimum and maximum of the four R-R intervals varied from each other by less than 30 milliseconds, the segment was classified as VT. Otherwise, the segment was classified as SVT. Thirty milliseconds was determined experimentally to be the optimal value in discriminating SVT and VT for the patient population under study.

RESULTS AND DISCUSSION

Acceleration

A number of novel acceleration techniques were described to perform correlation-based morphology assessment. By using all of the aforementioned acceleration techniques, only 70 multiplications and 720 additions are typically required per classified beat. Accelerations of more than 15X were measured experimentally, compared with an approach using the full computation of (1) for both registration and classification. The computation required to implement the proposed system is within the limits of presently available microcontrollers.

Classification

The rate-based and correlation-based system's performance in classifying segments of five consecutive beats was then tested. Using the data from eight patients exhibiting both VT and SVT, 387 5 beat segments of VT were excised and 944 segments of SVT were excised and classified by both techniques. The rate-based system achieved an error rate of 5.7% (22/387 misclassified) for VT segments and 5.5% (52/944 misclassified) for SVT segments. The correlation-based system had an error rate of 0% (0/387 misclassified) for VT segments and just 0.2% (2/944 misclassified) error rate for SVT segments. The results are summarized in Table1.

	Rate-Based System		***Correlation-Based System***	
	VT	**SVT**	**VT**	**SVT**
VT	365	22	387	0
SVT	52	892	2	942

Table 1 -Experimental Results

Given the small number of patients in the study, care must be taken in interpreting the results. Nevertheless, these initial results are encouraging. At present, a larger test database is being constructed to more definitively benchmark the proposed system.

CONCLUSION

A computationally-efficient, arrhythmia classification architecture using correlation-based morphology assessment was presented. The architecture classifies individual heart beats by assessing similarity between an incoming cardiac signal vector and a series of prestored class templates. A series of these beat classifications are used to make an overall rhythm assessment. The system makes use of several new results in the field of pattern recognition.

The resulting system achieved excellent accuracy in discriminating SVT and VT, and compared favorably to a rate-based approach.

Acknowledgments

This work was supported in part by a contract from the Office of Naval Research.

References

1. Myerberg R., Kessler, K., and Castellanos, A. "Sudden Cardiac Death: Structure, Function, and Time-Dependence of Risk", *Circulation*, Vol. 85, No. 1, January, 1992, I2-I10.
2. Singer, I., *Implantable Cardioverter Defibrillator*, Futura Publishing Co., Armonk, NY, 1994.
3. Wilkins, J., *Correlation Based Time Series Pattern Recognition for Implantable Defibrillators*, Ph.D. Thesis, Information Systems Laboratory, Stanford University, Stanford, CA, June, 1996.